REPORT

2026 Identity Security Landscape

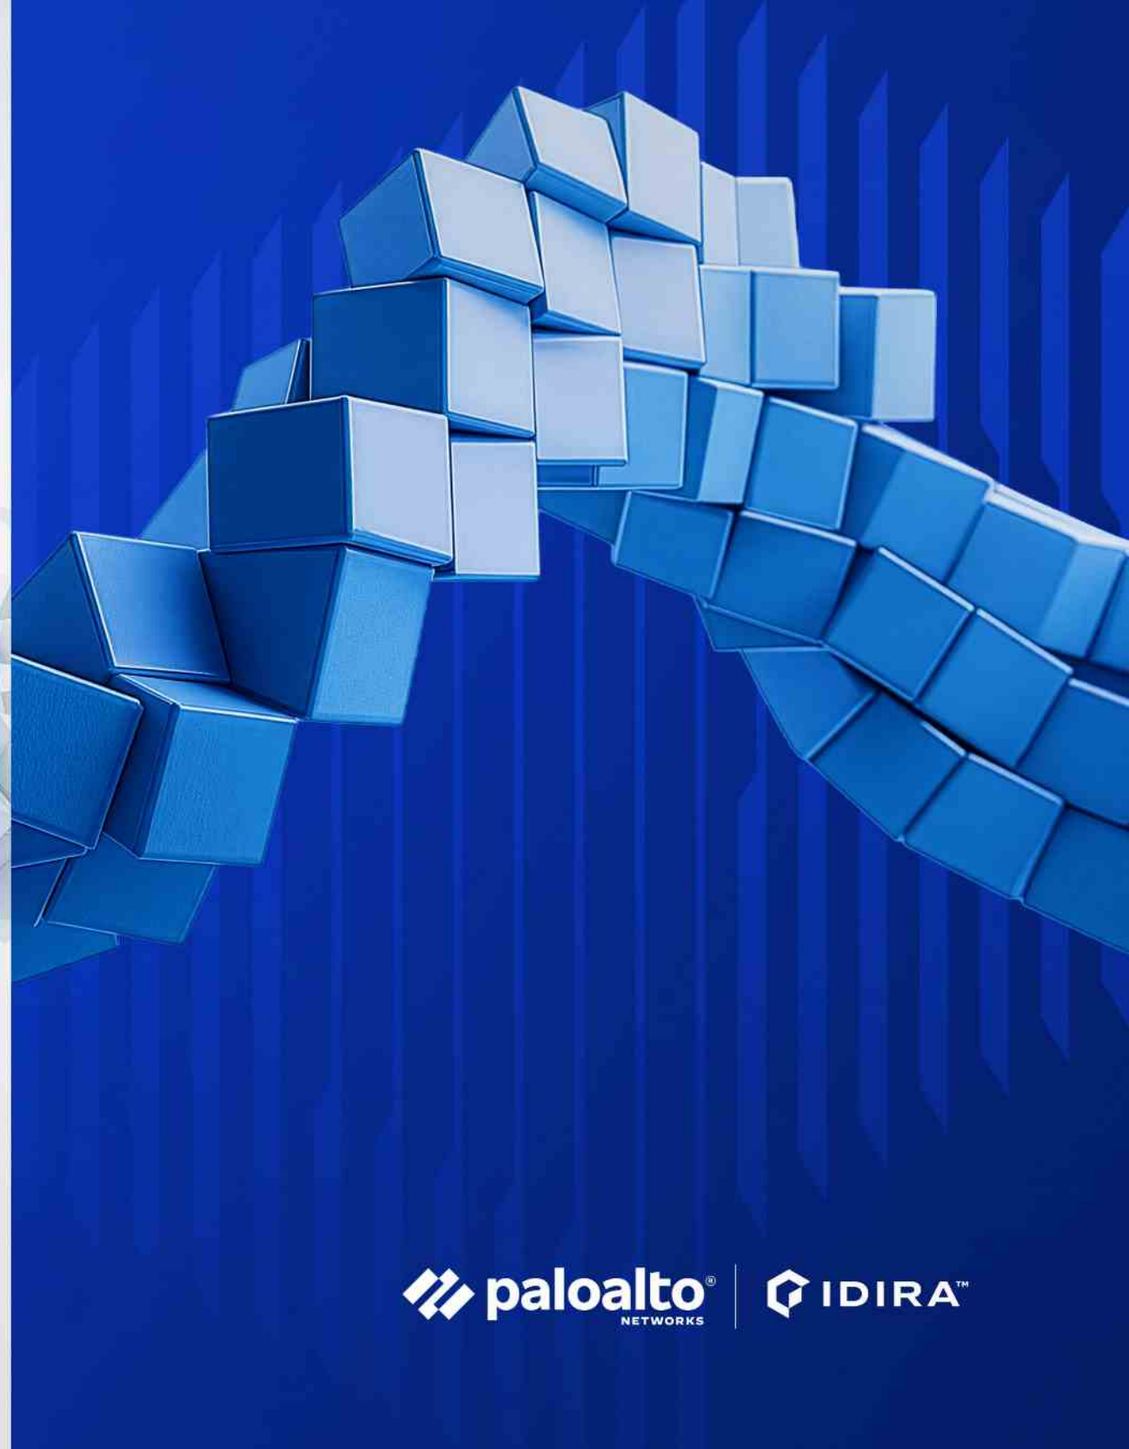

paloalto NETWORKS | IDIRA™

Table of Contents

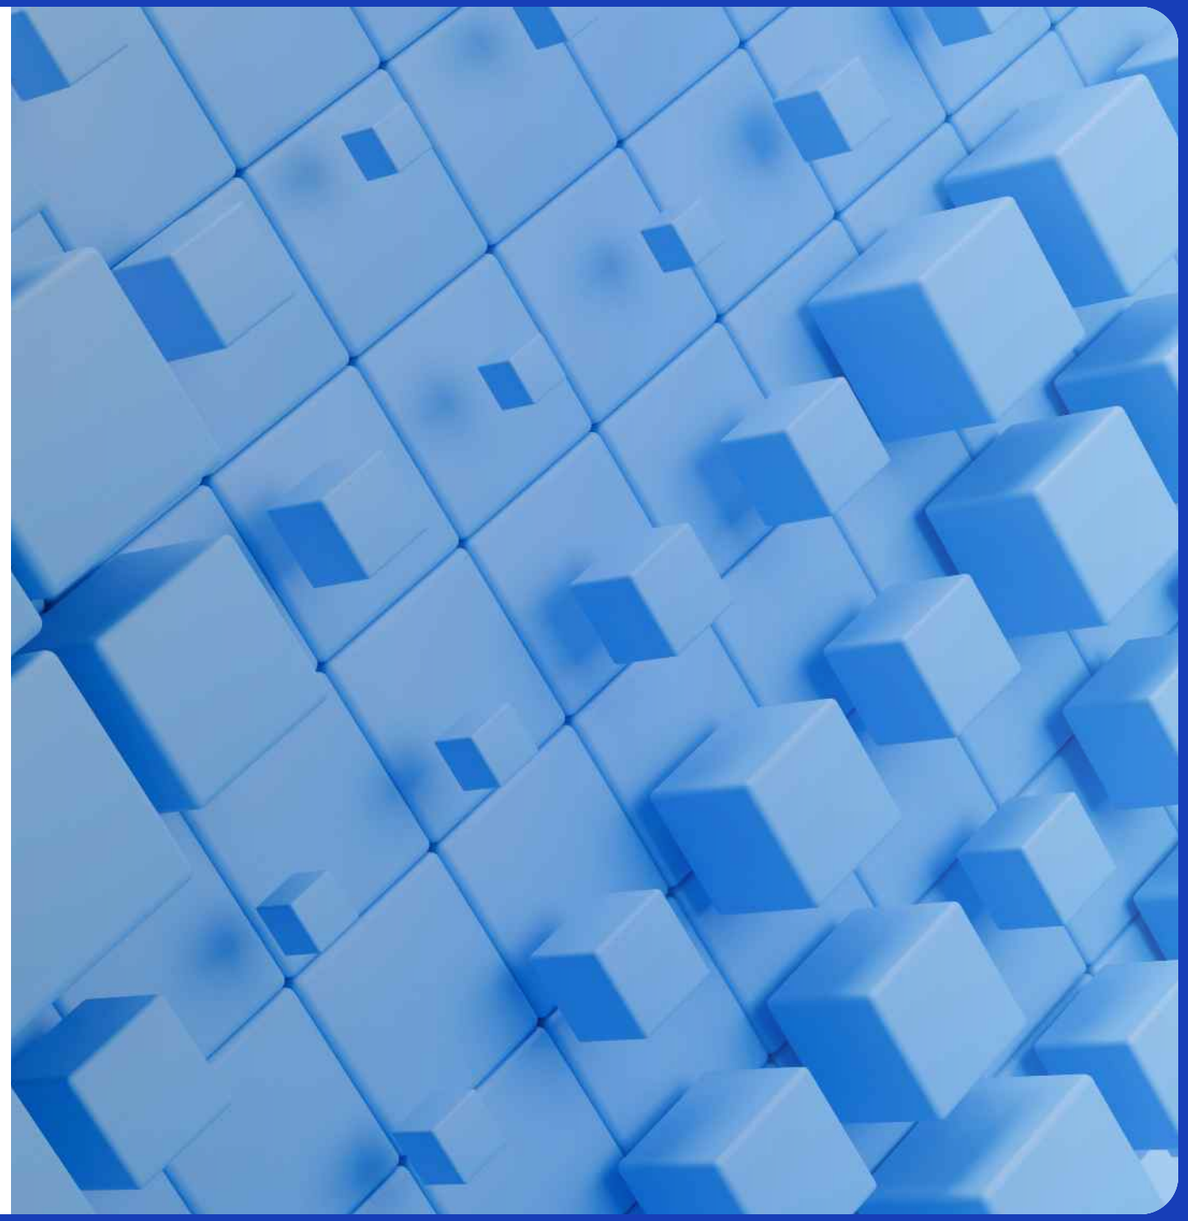

Executive Overview

Executive Overview

Welcome to the *2026 Identity Security Landscape* report. First, a sincere thank you to the 2,930 cybersecurity decision-makers worldwide whose generous insights made this research possible. And, congratulations to you, new and returning readers. You made it. In the data, we think you'll find a mirror for what you're seeing in your day to day—a world that's expanding faster than the tools we rely on to govern it.

Organizations are currently contending with an average of 109 machine identities for every human. AI agents are the fastest-growing subset of that population. We are now outnumbered in a world of service accounts, workload identities, and AI agents authenticating through certificates, tokens, keys, and secrets.

Our race to onboard the agentic workforce brings with it a new era that warps time and collapses trust. We were already dealing with global uncertainty, economic pressure, and nonstop transformation. Now adversaries are impersonating humans at scale with AI, and frontier models are beginning to surface vulnerabilities and mobilize exploit chains faster than defenders can patch. Our emotional center of gravity has shifted beyond password anxiety to "Can I trust what looks and sounds human?"

In this timeline, everything ages at machine speed. Trust has a shorter half-life while human privileges persist beyond their intended use. Certificates expire faster, while tokens live longer than they should. Secrets sprawl where they should not, while role assignments outlive the business that created them.

Machine identities now outnumber human identities by **109:1**.

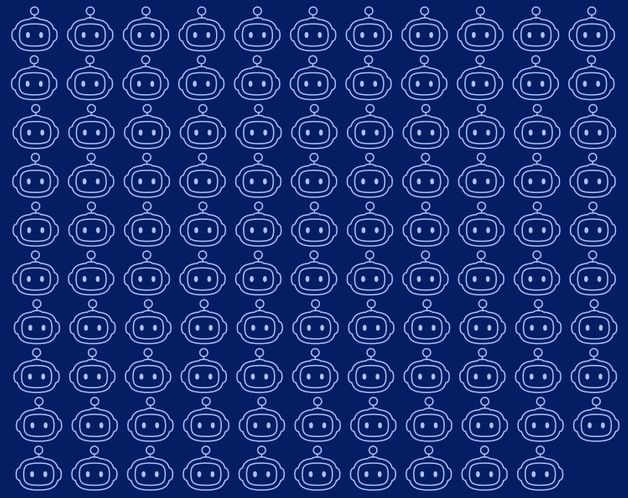

EXECUTIVE OVERVIEW

The 2026 data presents a defining risk: the speed gap between control intent and control execution. Organizations can define identity security principles in theory. In practice, however, real-world environments are so fragmented and dynamic that identity systems no longer function as designed. Nine out of 10 organizations experienced an identity-related breach in the last 12 months, with 83% reporting at least two such breaches.

Practitioners must battle AI-enabled adversaries with an identity stack that was fragmented long before agents arrived. They are moving at full speed in low visibility, trying to secure humans, privileged access, service accounts, secrets, certificates, and workload identities across disconnected control planes—none of which can intercept abuse at machine speed. It only takes one compromised access path to create cascading security and operational risk.

The old operating model that relies on human-centered identity architectures and static access tools is becoming administrative fiction. We see this shift as the catalyst for AI-driven platformization.

Good plans start with good data. Whether you've been tracking these trends for years or this is your first scouting mission, we hope this report provides you with a clear vantage point of the 2026 landscape.

Let's get to the numbers.

Amy Blackshaw
SVP, Identity Security
Palo Alto Networks

What's Inside

- **The species shift** of the machine-to-human identity ratio and why the economics of the 47-day public TLS certificate mandate is pushing certificate automation from nice to have to an operational requirement.
- **The 12-hour fragmentation tax** and how converged identity controls can reclaim analyst time and improve response.
- **How cyber insurance, compliance, and partner expectations** are guiding identity security investments.

CHAPTER 1

Identity Beyond Human Scale

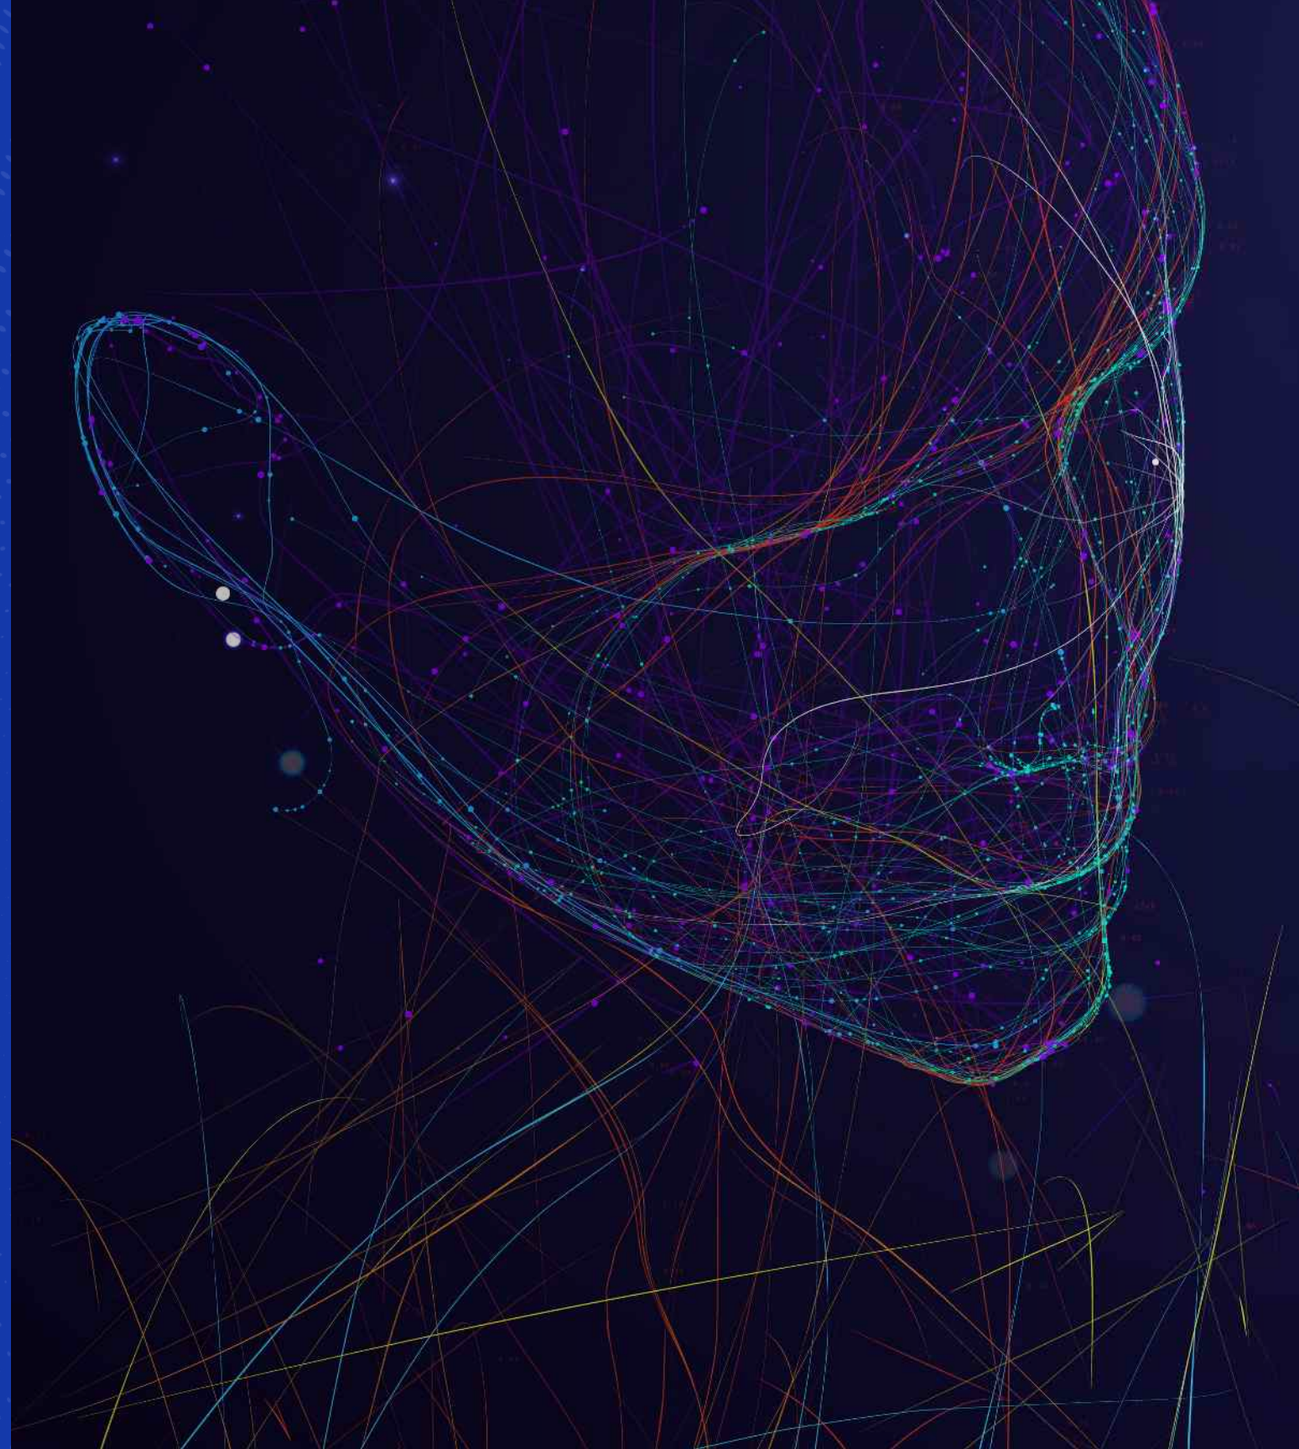

paloalto NETWORKS | IDIRA™

Identity Beyond Human Scale

Among the faces of your 2026 organization chart, there are a few million coworkers you can't see. They communicate at speeds you can't process and scales you can't track. Humans are no longer the dominant species of our network. We are but a tiny minority in an ecosystem dominated by autonomous agents and machine workloads.

This report looks at how three distinct forces are colliding:

Humans: This flesh-and-blood workforce is now vastly outnumbered but carries more privileges than ever before. Today, every employee is a high-value target that must be shielded from new risks.

Machine identities: These identities now outnumber humans by 109 to 1, spanning service accounts, secrets, certificates, tokens, workload identities, and now AI agents. They are becoming harder to discover, rotate, and govern.

AI agents: These software actors can reason, act, call tools, and increasingly orchestrate task agents with minimal human intervention. This class of privileged machine identity has created a governance vacuum that only AI-driven controls can fill.

This year's survey found that AI and LLMs are now tied with standard machine identities as the primary driver of identity growth (figure 1). Across environments, the latest ratio (109:1) of machine identities to humans includes 79 AI agent identities per human.

WE ASKED | What are the main factors driving the increase in identities in your organization over the next 12 months?

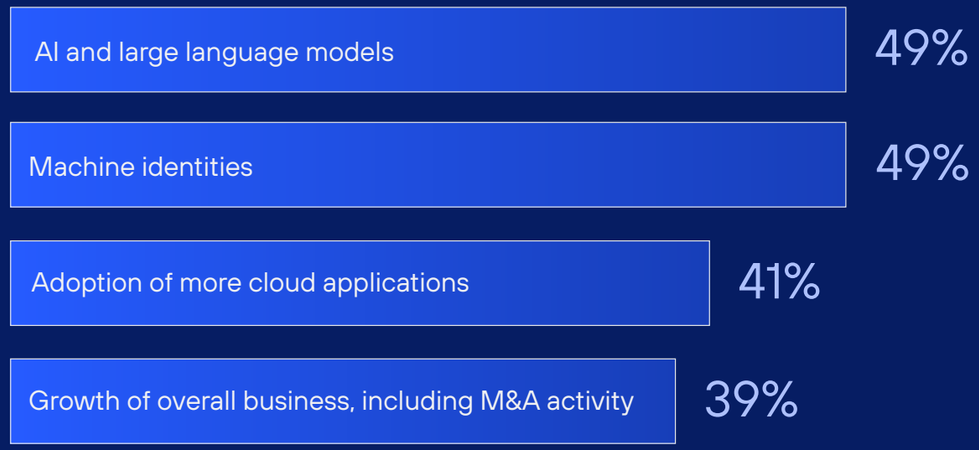

Figure 1. AI and machine identities are now the joint top drivers of identity growth (n=2,844)

AI agents represent the fastest-growing subset of this population. Not every machine identity is suddenly agentic, but autonomous behavior is being layered onto an already dense and undergoverned stack. Identity has shifted from a provisioning problem to a runtime control problem.

Over the next 12 months, organizations expect AI agents to increase by 85% and machine identities to increase by 77%, compared to the 56% growth in human identities. Moving forward, identity volume is no longer a mirror of organizational headcount. Every human hire now brings a swarm of nonhuman identities with them. And, every new application, workflow, or AI-enabled process introduces additional access paths.

The Myth of Bounded AI

"Every identity has a purpose, but few have a plan" might best describe the modern enterprise. This year, 99% of organizations report using AI agents. Some summarize meetings, and some write or refactor code. Others invoke external tools or data sources. Most organizations can explain why these agents were created. Far fewer can say, with precision, what they can access, how that access is bounded, how it's revoked when conditions change, and who or what can inherit that access.

Survey respondents indicate that, on average, 40% of AI agents and 40% of machine identities already have access to organizational data. This data might include, for example, financial records, customer personally identifiable information (PII) or protected health information (PHI), intellectual property, critical operational technology, and high-value business systems.

At the same time, 90% of respondents agree that AI agents should operate under least privilege principles, with bounded access and tighter controls, and 91% believe they can rapidly contain compromised AI agents. Yet, we aren't seeing evidence of consistent controls capable of stopping them if and when agents deviate from their intent.

of organizations agree that AI agents should operate under the principle of least privilege.

Identity Security in Theory Doesn't Always Work in Practice

Containment exists more on paper than in practice. Across AI agent types, fewer than half of organizations report applying core lifecycle controls, such as behavioral monitoring and credential revocation (figure 2), with the ability to constrain or shut down agent activity remaining inconsistent.

These controls matter because AI agents don't behave like conventional users. They don't pause or double-check the access they have been given. Instead, they act on that access at machine speed.

The broader environment is already struggling with this problem. On average, only 39% of privileged access is managed through a just-in-time (JIT) or zero standing privilege (ZSP) model. Most organizations still rely on standing access. Adding AI agents, workloads, and connectors into these environments explodes the number of identities that security teams must control correctly and continuously. Deployment that outpaces oversight is how a governance vacuum forms.

Many organizations can identify where AI agents exist. Visibility matters, but visibility without consistent discovery, policy enforcement, and real-time control is just a dashboard. Without this scaffolding, security teams lack the critical kill switch that determines whether their organization can safely adopt AI at scale.

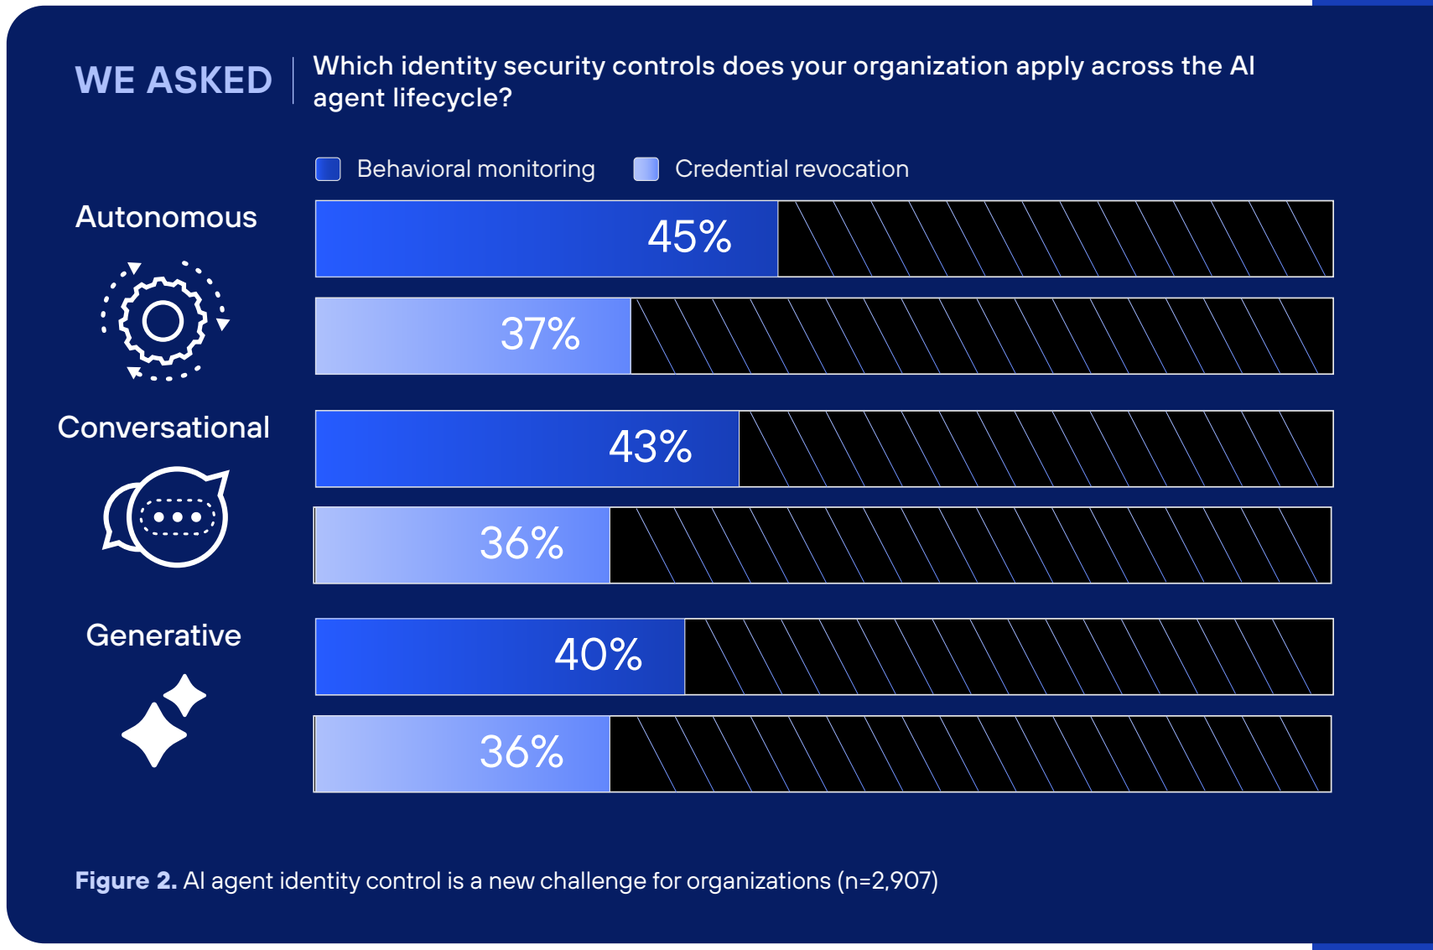

Figure 2. AI agent identity control is a new challenge for organizations (n=2,907)

The Identity Management Gap

The data points to an operational fault line between leadership perception and practitioner experience. Among C-suite respondents, 54% believe their organization is completely effective at continuously enforcing least privilege, while 61% of the practitioners doing the work disagree. Both can sound right, depending on whether you're looking at a summary or the actual logs.

This disconnect is the operational expression of the "speed gap." It starts with how the enterprise defines the problem. Most organizations still view privilege through a human lens, leaving out the nonhuman identities that now do the bulk of the work (figure 3).

Closing this gap requires granular control. Modern containment means real-time orchestration across identities, sessions, and environments. Today, most organizations don't have that capability. Instead, to keep services running, organizations are leaving broad permissions in place for machines and compensate by overtrusting the humans who supervise or invoke them.

Power Overwhelming: The Risk of Uncontrolled Privilege

The paradox of modern IT is that, as humans become the minority identity, our individual risk profile and individual power have skyrocketed. A single human identity now represents a user and a control point.

Today, one login sits at the center of a large collection of downstream access, approvals, integrations, and machine activity. This center of control has expanded dramatically, where a single action can trigger workflows, invoke agents, access sensitive systems, or move data across environments. It has turned a standard login into a one-stop shop for extortion groups who know that most identity controls weaken after authentication and plan accordingly.

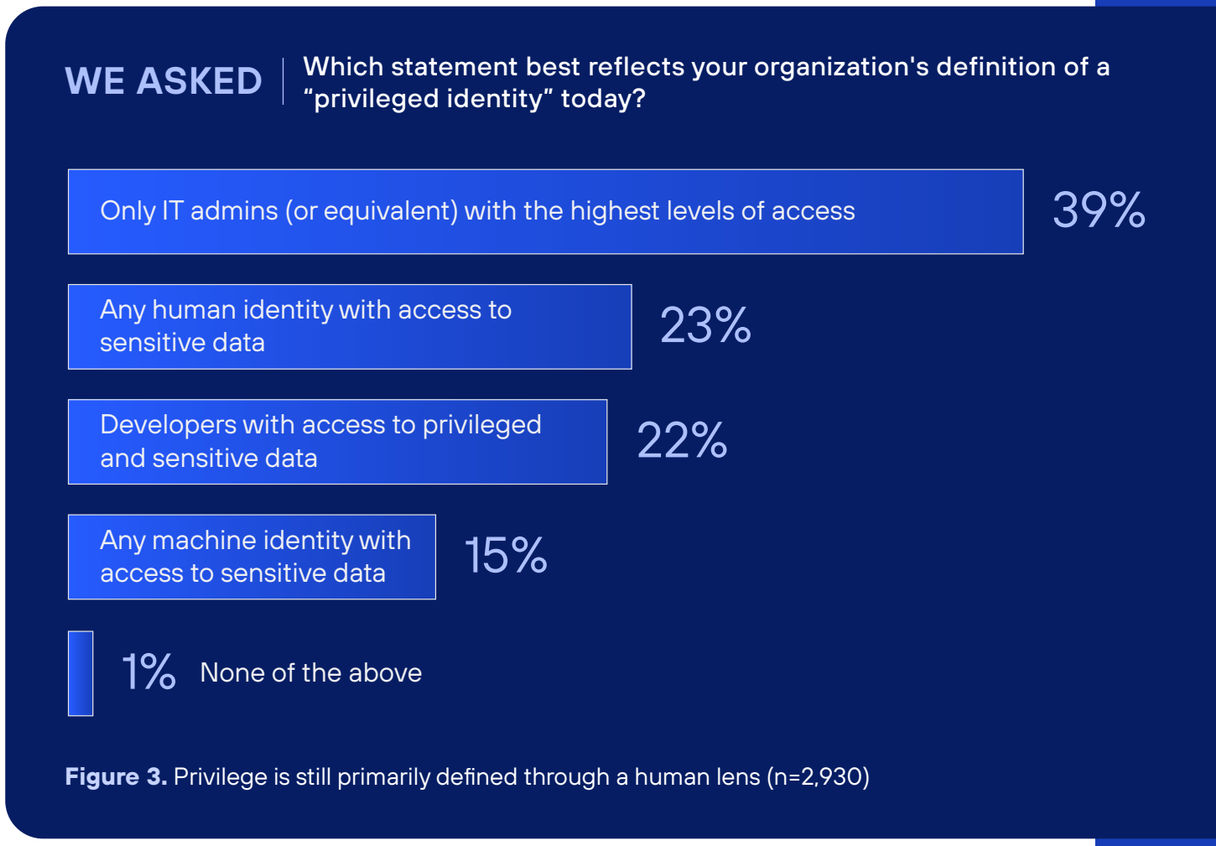

Figure 3. Privilege is still primarily defined through a human lens (n=2,930)

Admin Rights for Everyone

The data is categorical in that 96% of respondents report that human identities operate with access far beyond what is required for their roles. And, on average, 42% of the human workforce has direct access to organizational data (figure 4).

Endpoints have also become more powerful. The modern workstation has evolved from a passive client into an **agentic endpoint**—a high-throughput execution environment where packages, browser extensions, AI tools, coding agents, and scripts are installed and allowed to act with startling speed.

This shift raises the value of endpoint least privilege. Many identity incidents become material because the endpoint still allows local admin rights or ungoverned process elevation, creating openings for credential dumping and browser token theft. Endpoint least privilege reduces the number of human identities that can turn a compromised session into lateral movement or outright data access.

The tectonic pressure of the 109:1 ratio, a fragmented control stack across access, privilege, endpoint, and machine identities, and the need to keep systems running has pushed organizations to the same compromise: broad access first, cleanup later. Velocity is preserved, but uncontrolled privilege is normalized.

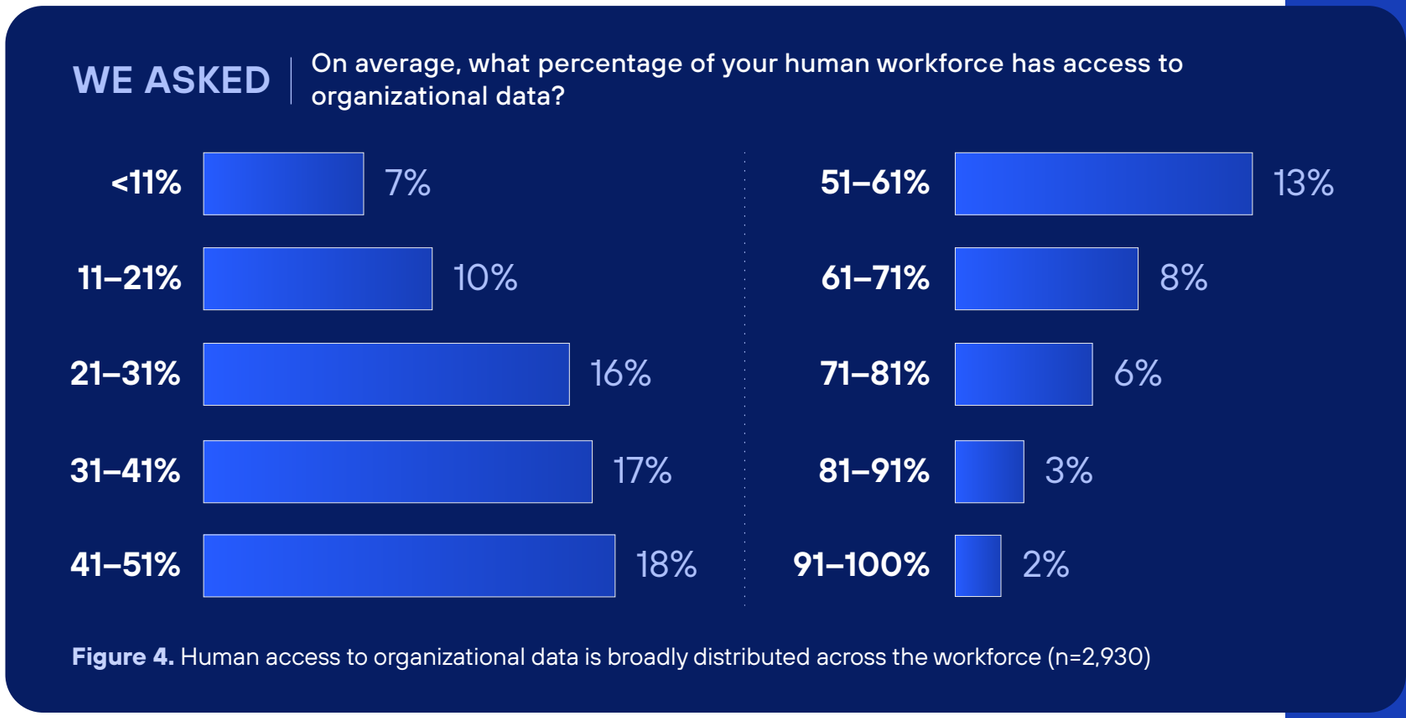

Figure 4. Human access to organizational data is broadly distributed across the workforce (n=2,930)

CHAPTER 2

Stop, Drop, and Wait:

The Identity Bottleneck

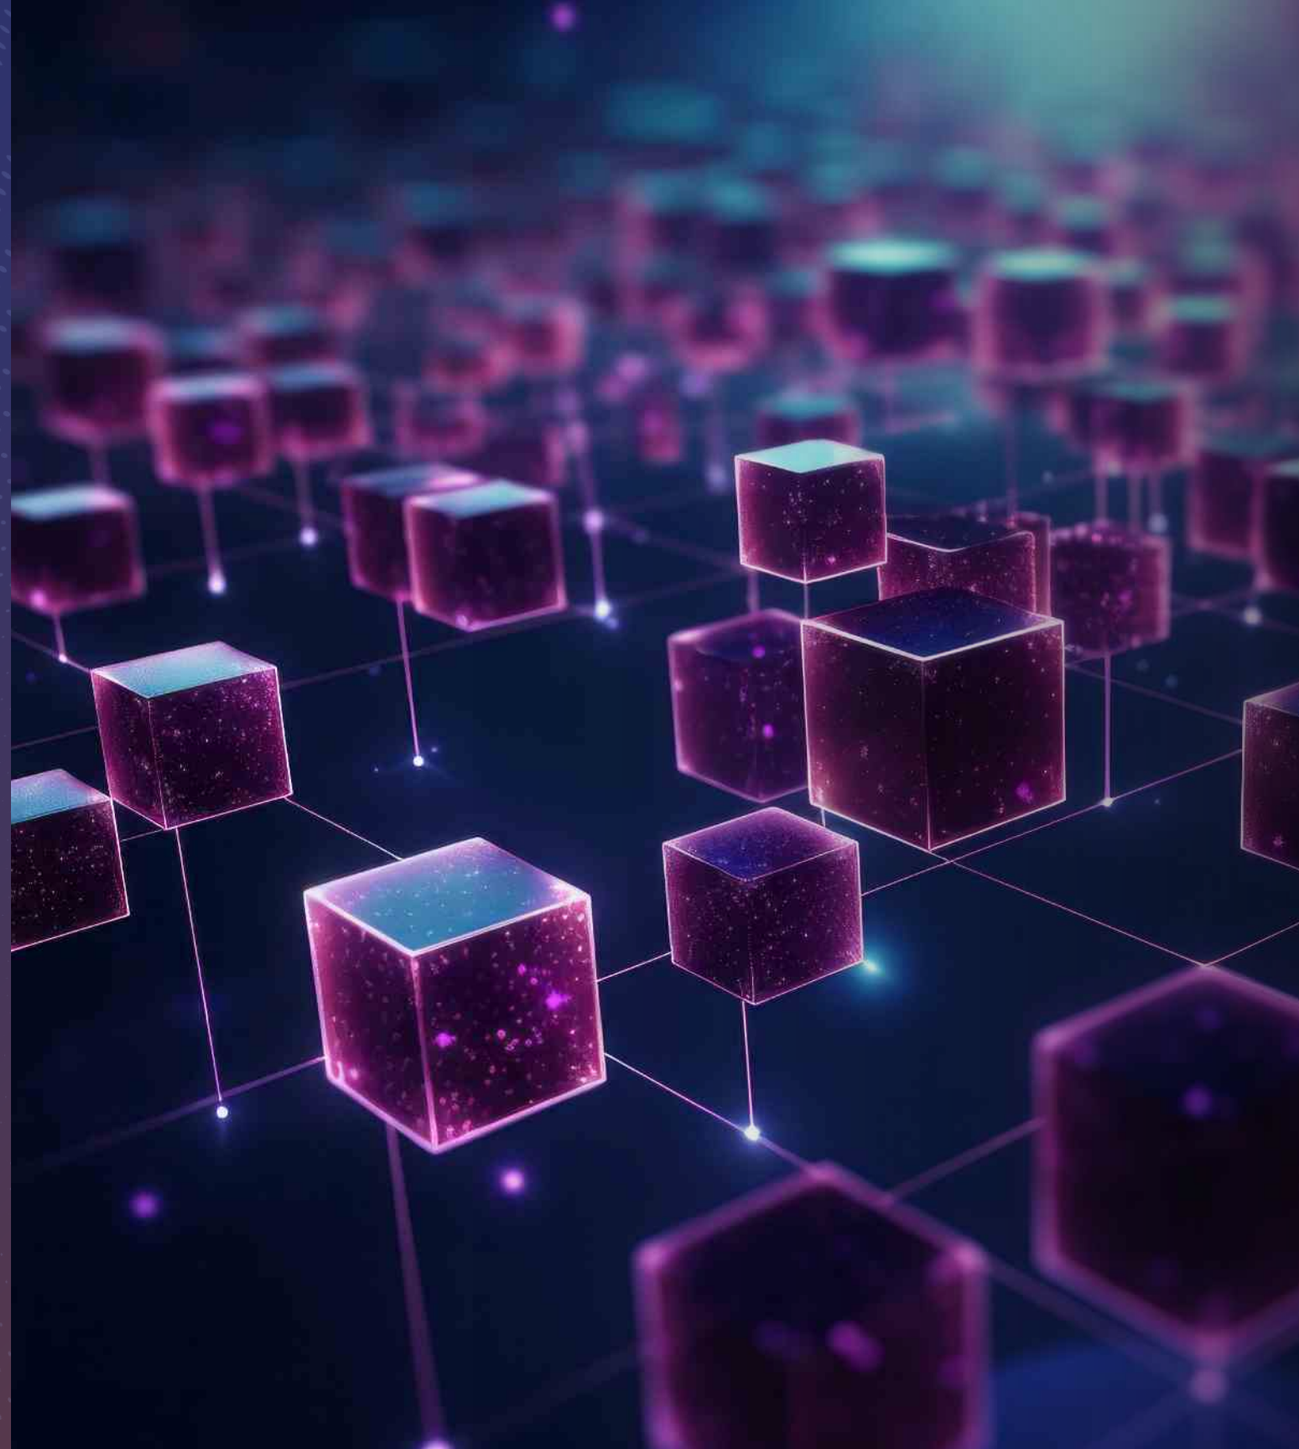

paloalto NETWORKS | IDIRA™

Stop, Drop, and Wait: The Identity Bottleneck

When an identity-related breach hits a fragmented environment, security teams are forced to correlate and respond across multiple consoles with incomplete context. Against AI-assisted workflows where attackers can run recon and access attempts across hundreds of targets in parallel, remediation feels like bringing a fax machine to a firefight.

The scale of this exposure is substantial. Across multiple identity-related breach categories, 90% of organizations report a successful identity-related breach in the last 12 months, with 83% seeing it happen at least twice, and 76% experiencing it three or more times.

Once an incident begins, many organizations struggle to assemble a coherent view quickly enough to contain it. Across more than 750 major cyber incidents investigated in 2025, Unit 42® found that 87% required evidence from two or more distinct sources to establish what happened. In complex cases, investigators drew on as many as 10 sources.[1]

1. *Unit 42 Global Incident Response Report 2026*, Palo Alto Networks, February 17, 2026.

IN THE LAST 12 MONTHS:

90%
report at least one successful identity-related breach

83%
of organizations experienced at least two successful identity breaches in the last 12 months.

76%
of organizations experienced at least three successful identity breaches in the last 12 months.

The Fragmentation Tax

This operational friction carries a heavy and measurable cost. In our research, 85% of organizations agree that fragmented identity systems and tools impact or delay their ability to detect and respond to identity-related threats.

The findings are more stark in the trenches with 97% of respondents reporting that tool fragmentation adds additional time to identity-related incidents, amounting to an average of 12 hours per incident. Meanwhile, the fastest real-world intrusions now reach exfiltration in just 72 minutes.[2] This fragmentation tax is the mathematical result of trying to find a needle in a pile of needles.

Practitioners don't lack the will to act. They are subsumed by scale and less sure about what they are looking at. Is the activity legitimate or an impersonation? Or, is it hijacked, inherited, or drifted? Security is battling a speed gap. The business moves at machine speed, but control does not.

While over half of respondents indicate they respond with password resets, fewer can rapidly remove privilege, revoke tokens, or terminate live sessions with the speed required to contain machine-driven abuse (figure 5).

WE ASKED Which identity-based response actions are included in your organization's incident response procedures?

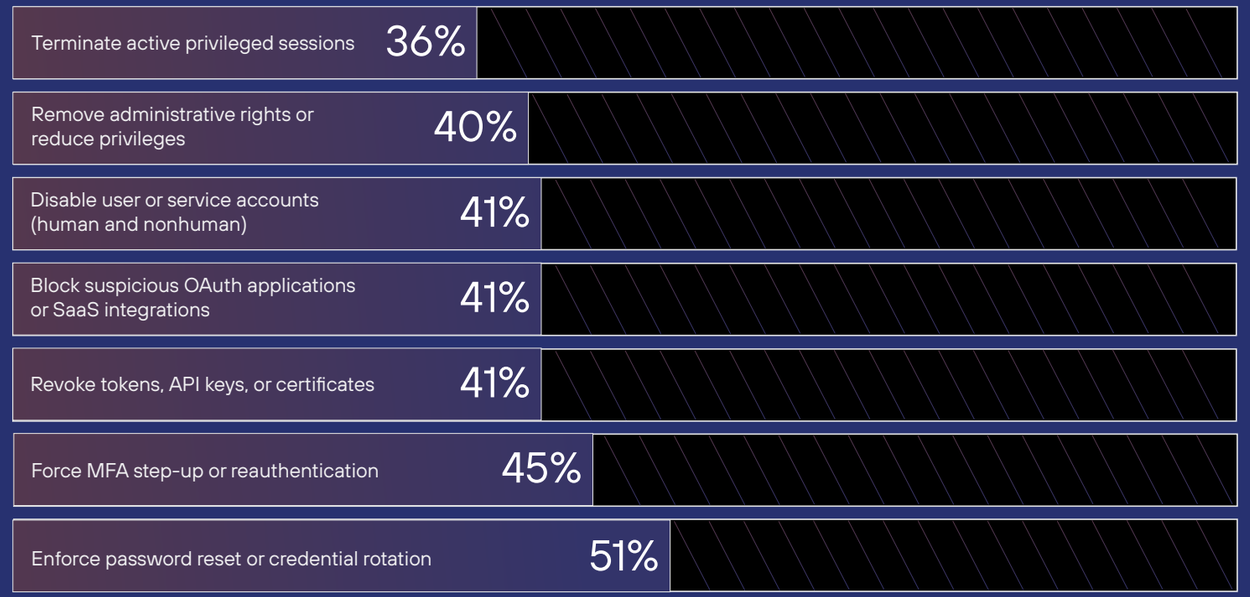

Figure 5. Most organizations still rely on basic identity response actions during an incident (n=2,930)

2. Palo Alto Networks Unit 42, *Global Incident Response Report* 2026.

Autopsy of the Authentication Model

Reliance on basic password resets while struggling to revoke live access exposes a critical architectural flaw. Many environments treat authentication as the primary control point with limited protection beyond the login.

For years, companies have relied on service accounts and machine identities to broker trusted access. AI-powered "brains" are now being grafted onto these identities, dramatically expanding their autonomy without a corresponding shift in control. This complexity is why 84% of respondents believe their organization could at least moderately improve awareness of the permissions and access granted to connectors and service accounts.

Single sign-on (SSO) and multifactor authentication (MFA) are necessary, but they are front-door defenses, not a complete identity security strategy. Once a user, token, connector, or agent is inside, the critical question is: What can it reach, approve, or leave behind? Currently, 56% of organizations report that they can't effectively enforce continuous least privilege for service accounts across cloud, SaaS, and on-premises environments. In the agentic enterprise, authorization must move closer to the action itself.

The Ghosts of Access Past

Security teams know exactly where the bleeding starts—with stale, standing access. One respondent mentioned "disgruntled workers or former employees that still have access to things they shouldn't." Others pointed to the quiet accumulation of risk, warning that "login credentials often remain active for years after a project has been completed" and that "elevated permissions often remain unchanged" long after a user's role evolves. Respondents also pointed directly to Joiner-Mover-Leaver (JML) pain, warning that "permissions from sister companies remain in legacy systems," creating "undocumented and unnoticed security holes."

This grant-now, fix-later mentality extends to third-party access and supplier integrations that are built on a shared-trust model, where "fix later" rarely happens: assess the vendor, review the access, collect the paperwork, and move on with your day. When a credential is abused, however, organizations need containment, not better documentation. The same logic applies to federated access and SSO. Once a user or partner is through the front door, what rules can they rewrite? Can they leave the backdoor unlatched for future use?

The market's answer to this exposure is usually consolidation, but vendor consolidation often reduces logos without reducing fragmentation. True platformization is a single, integrated codebase that eliminates blind spots, not a disconnected tech stack.

84%

of respondents believe their organization **could at least moderately improve its awareness of the permissions and access** granted to connectors and service accounts.

CHAPTER 3

The New Half-Life of Trust

paloalto NETWORKS | IDIRA™

The New Half-Life of Trust

In 2026, identity is a ceaseless, tumultuous stream of requests, tokens, sessions, certificates, and machine actions that must be filtered and verified. Our data suggests that practitioners are tracking the speed gap and backing away from the illusion of static trust and check-the-gate defenses. A new urgency exists for continuous, real-time authorization and granular identity detection and response (figure 6).

Seeing Is No Longer Believing

What makes the current trust problem different from older phishing eras is that attackers are assembling believable people and believable activity, not merely stealing credentials. We are seeing the rise of identity stitching, where attackers use AI to mine open-source intelligence from social media and corporate directories to sew together convincing synthetic profiles. As these stitched identities continue to scale the uncanny valley, social engineering becomes increasingly convincing and leads to hard-to-detect access that looks like it belongs.

This part is one of a broader trust collapse. Many environments are filled with hard-coded secrets that AI-assisted code search can surface instantly, OAuth tokens that outlive their business purpose, certificates that require faster renewal, and machine credentials across sprawling automation. If any of these artifacts are overexposed or overtrusted, the window of misuse can outlast the context that created them.

WE ASKED | Over the next 12 months, what areas are seen as priorities for your organization's identity security program?

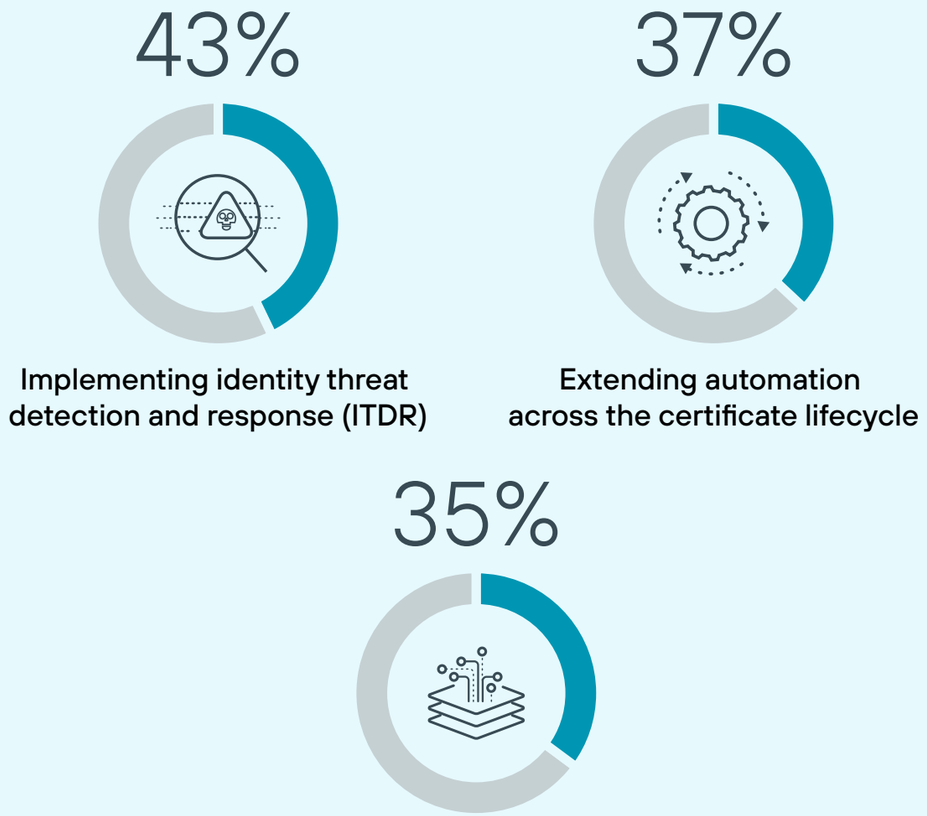

Figure 6. ITDR emerges as a top identity security priority for the next 12 months (n=2,930)

Schrödinger's Breach

Is your encrypted property perfectly secure, or is it already breached? In the harvest now, decrypt later (HNDL) era, the answer is both. It is the quantum superposition of modern security. Data stolen today sits in a state of pending breach until a cryptographically relevant quantum computer (CRQC) finally opens the box. Awareness of this threat is high, yet 64% of organizations still report they are not yet prepared to defend against quantum-enabled threats, while a further 86% agree that they will need external help to become quantum ready.

The 48,000-Hour Wake-Up Call

While quantum pressure builds on the horizon, certificate operations are already testing business continuity. Public TLS certificate lifetimes are stepping down this year from 398 days to 200, then 100, and soon 47. While this reduction doesn't define every machine certificate, it adds to a growing set of operations that break down under manual processes and limited visibility at enterprise scale.

Tracking the lifespans of thousands of public TLS certificates already represents a major operational burden. Behind this sits a much larger estate of internal certificates, keys, secrets, and algorithms. Managing and securing this hidden infrastructure requires centralized visibility, inline automation, and crypto agility.

Despite this deadline, 71% of organizations don't fully automate certificate renewal and monitoring across all environments (figure 7), and 98% report challenges securing PKI (figure 8).

WE ASKED | How prepared is your organization for the upcoming 47-day mandate?

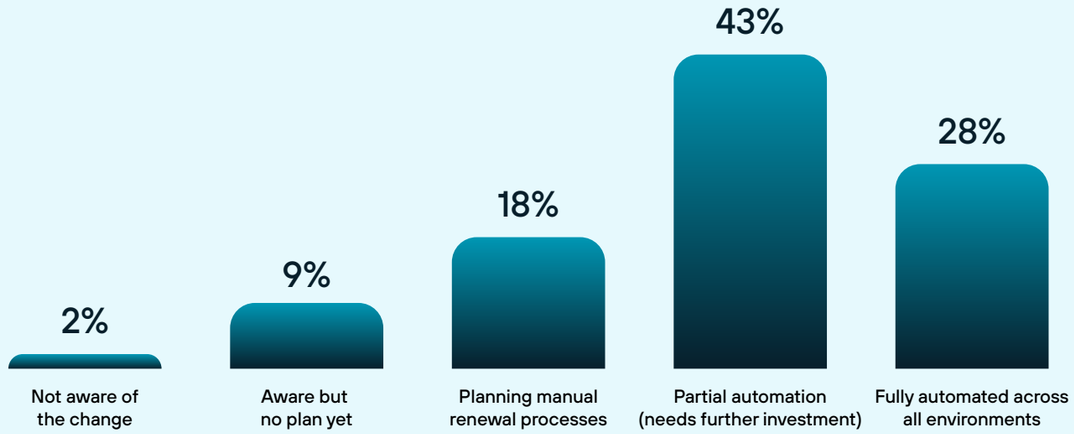

Figure 7. Most organizations have not fully automated certificate operations (n=2,930)

WE ASKED | What are the primary challenges or barriers your organization faces when securing public key infrastructure (PKI) and certificates?

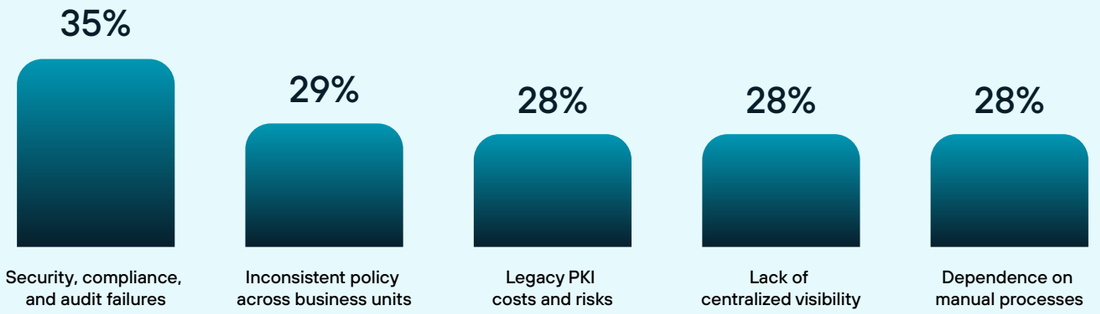

Figure 8. Policy gaps, visibility gaps, and manual processes constrain PKI and certificate security (n=2,930)

Death by a Thousand Renewals

To put the TLS certificate mandate in perspective, let's do some quick math. The transition to 47-day lifespans triggers an eight- to twelve-fold increase in manual workloads. For an organization managing just 1,000 certificates, respondents estimate an average financial impact of $272,088 from certificate-related failures. And, the operational overhead would skyrocket from roughly 4,000 hours to 48,000 hours.

If organizations can't automate certificate rotation now, they will carry that same operational burden into the crypto-agility demands of post-quantum transitions later.

Identity Is the New Credit Score

In 2026, identity security directly affects revenue, risk, and viability. Frameworks, like NIS2 and DORA, dictate that an organization's identity posture now determines everything from regulatory standing and partnership potential to its insurability.

Advanced identity controls have become a strict prerequisite for coverage, with insurance mandates now shaping strategy for nearly every organization. Among respondents, 98% confirm that insurance requirements influenced their identity security investments over the last 12 months, with 81% reporting that this influence was moderate or significant (figure 9).

WE ASKED | Which factors most influence your organization's decision to implement identity security controls?

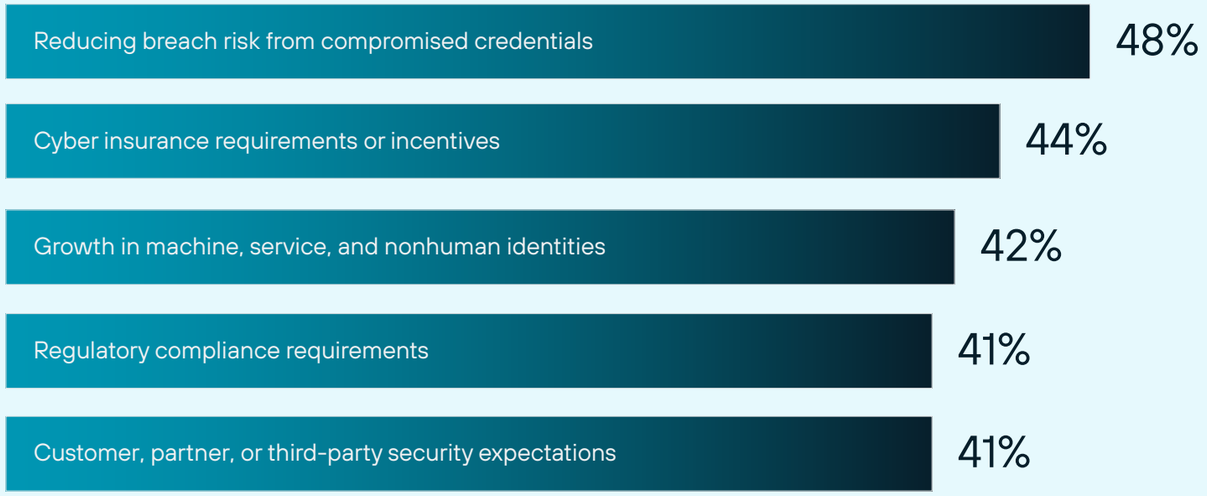

Figure 9. Breach risk and insurance pressures are the top drivers of an identity security investment (n=2,930)

WE ASKED | Looking ahead to 2026, what identity-related risk concerns you most and why?

Figure 10. AI is the top identity concern for 2026 (n=2,930)

A New Time Compression: Weaponized Intelligence

We are entering a new asymmetry between machine-speed offense and human-speed remediation. Frontier models, notably Anthropic's Claude Mythos Preview, are shrinking the time required to find vulnerabilities, understand exploit paths, and draft usable attack logic. Mythos Preview has already identified thousands of zero-day vulnerabilities across every major operating system and web browser.

Under this kind of patch stress, every exposed secret, misconfiguration, or standing admin path becomes an immediate liability. When code cannot be fixed fast enough, identity becomes the critical control plane that can still adapt in real time. It requires teams to strip away standing privilege, map hidden access paths, and enforce just-in-time access.

The data suggests that security leaders are pivoting toward this machine-speed reality. When we asked respondents to look past the architecture and name their top identity concern in 2026, 55% pointed to AI-enabled threats, including autonomous attack tooling and AI misuse. Phishing and deepfakes remain top of mind, along with machine identity and comprehensive access and privilege governance. The common thread across these responses is whether an identity, once active, can be trusted from one action to the next.

Recommendations

A Path to Centralizing Identity Security

A Path to Centralizing Identity Security

Bring the Agentic Workforce into the Light

Discover and centrally manage agents across SaaS, cloud, and developer environments. Enforce least privilege by granting agents access only for the duration of a specific task. Govern and audit agent actions to support compliance and understand what actions were performed by what agent and on behalf of which user.

Implement JIT and Time-Bound Enforcement Wherever Possible

Enforce least privilege across every access path. Eliminate standing access with just-in-time controls and automated rotation. Grant ephemeral access, revoke at session close, and audit every action.

Collapse Tool Silos and Automate the Identity Lifecycle

Reclaim the fragmentation tax, cut through operational mud, and govern at the speed of the systems that need protecting. In our time-compressed world of recurring breach and agentic execution, the organizations that win will be the ones that can move from signal to decision to action without crossing five consoles to do it.

Conclusion

From Entropy to Automation

paloalto NETWORKS | IDIRA™

From Entropy to Automation

This year's findings point to a simple reality. The technologies driving scale, speed, and efficiency are reshaping the identity problem faster than most organizations can govern it. The mission for security leaders is to verify a login. It is also to prove, in real time, that the requester is real, the access makes sense, the chain of trust has not drifted, and uncontrolled privilege is not sitting around waiting to be exploited.

When humans, workloads, tokens, service principals, and agent teams can initiate meaningful downstream actions, identity can't be relegated to an administrative function. Meeting this moment requires identity practitioners to embrace a next-generation control model capable of eliminating fragmentation and closing the loop between detection and enforcement.

We've scouted the terrain. Now we're rooting for the future you're working to build and the people you need to protect.

Let's build it together.

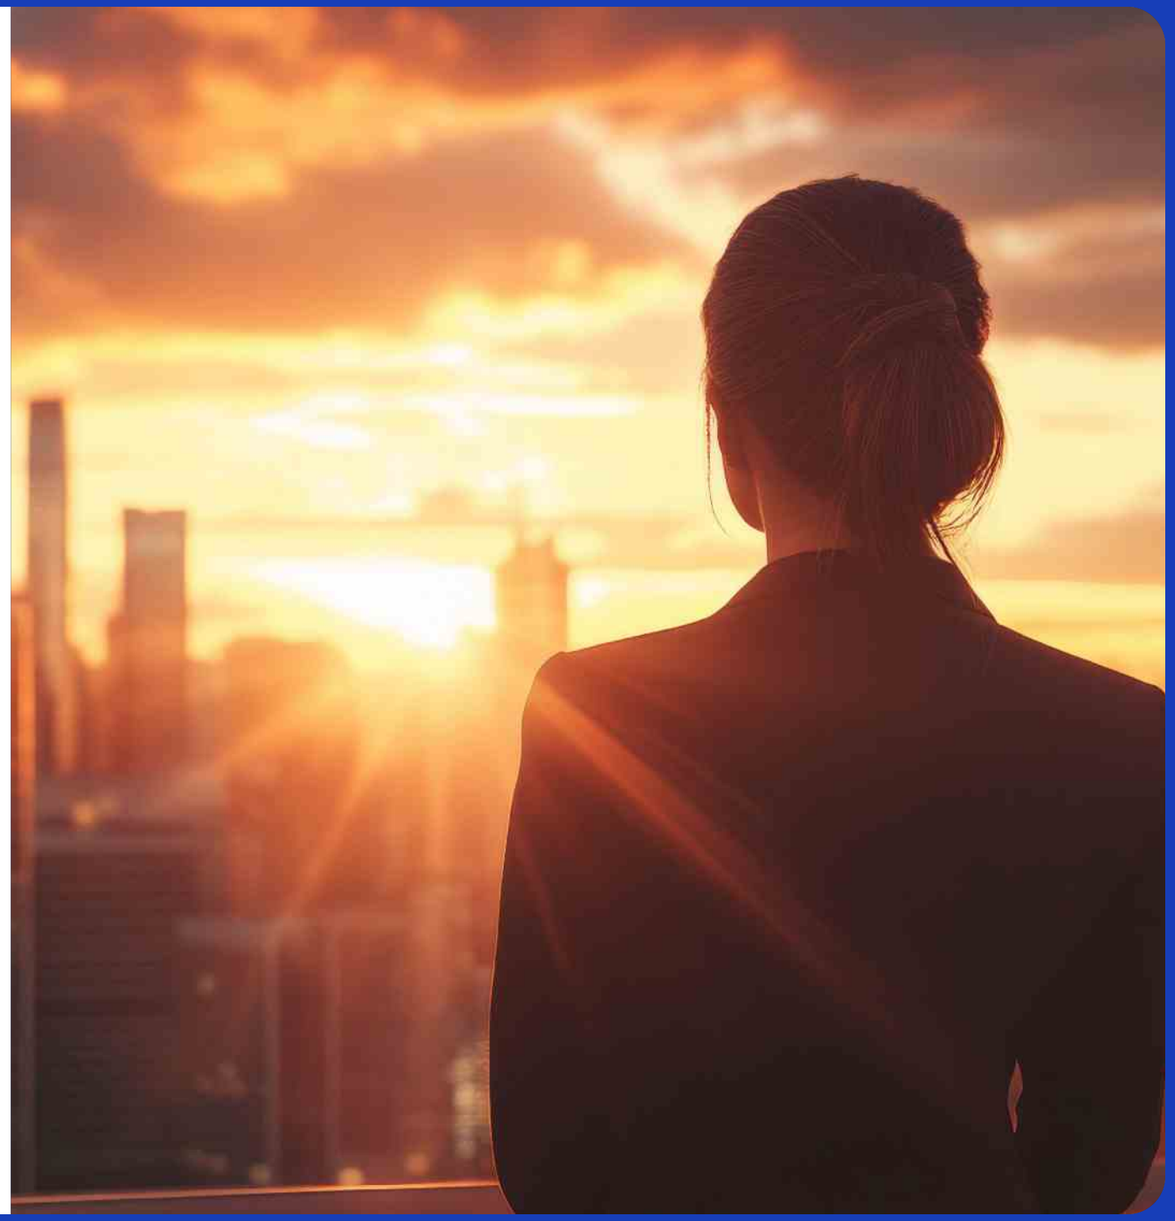

Research Methodology and Demographics

paloalto NETWORKS | IDIRA™

Research Methodology and Demographics

The 2026 Identity Security Landscape study, commissioned by Palo Alto Networks, was conducted by B2B technology research partner Vanson Bourne. The results of the study cited in this report were fielded across private and public sector organizations between March and April 2026. The sample included 2,930 cybersecurity decision makers based in Australia, Brazil, Canada, France, Germany, Hong Kong, India, Israel, Italy, Japan, Mexico, the Netherlands, Saudi Arabia, Singapore, South Africa, Spain, Taiwan, UAE, the UK, and the US.

The following graphs show the demography breakdown of the study's respondents.

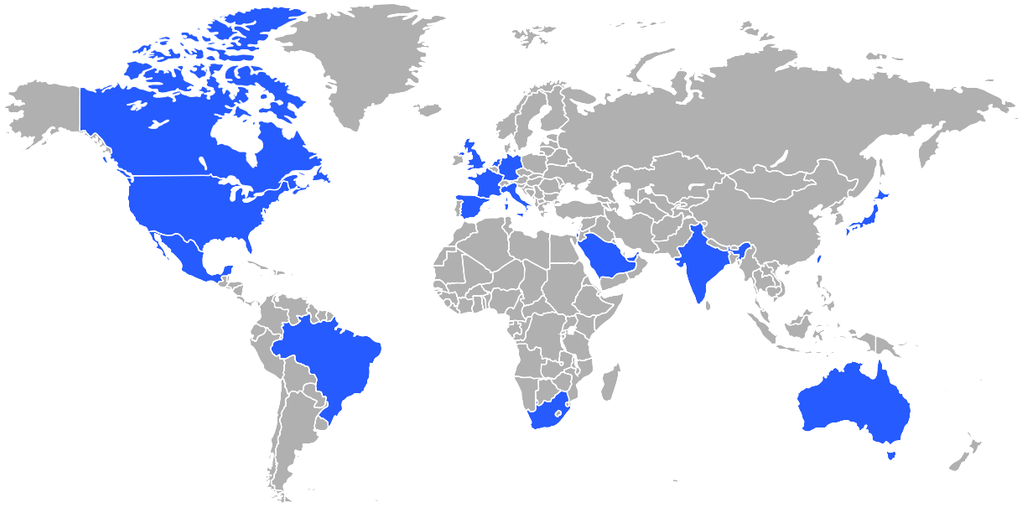

RESPONDENTS BY GEOGRAPHY

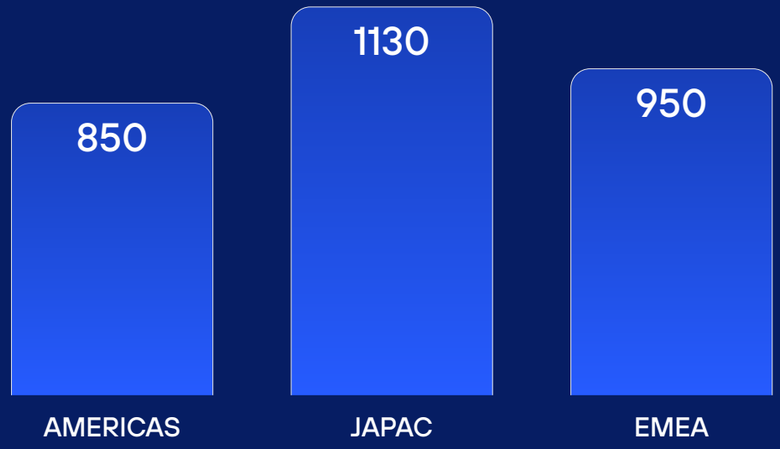

RESPONDENTS BY SECTOR

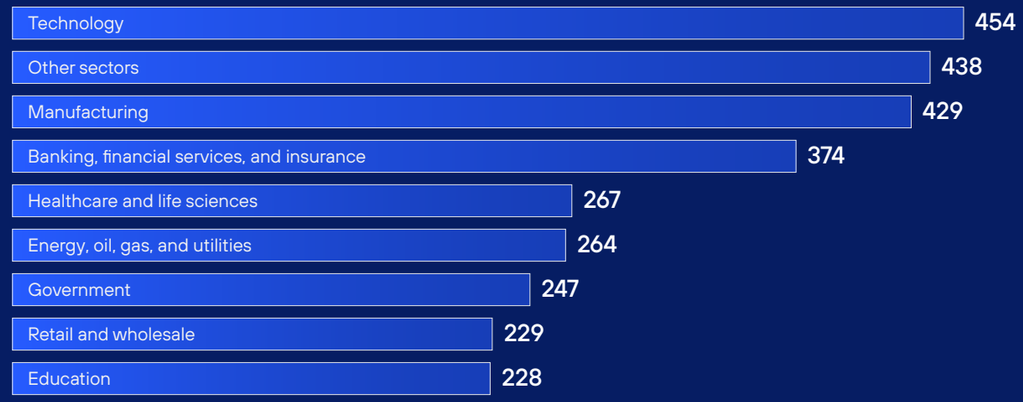

RESPONDENTS BY COMPANY REVENUE

SMB – $10.1 to $350 million USD	750
Core – Between $350 million and $1.5 billion USD	1,451
Majors – Over $1.5 billion USD	729

RESPONDENTS BY DEPARTMENT

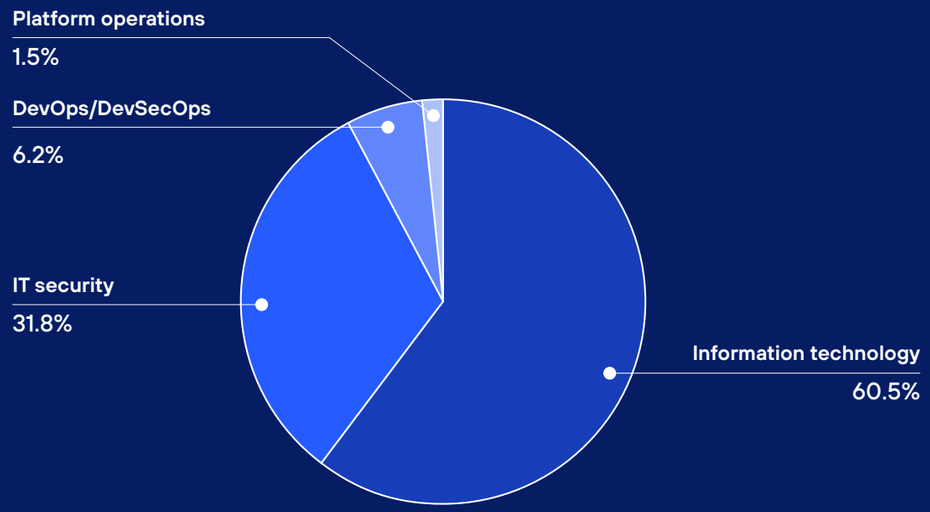

RESPONDENTS BY TITLE

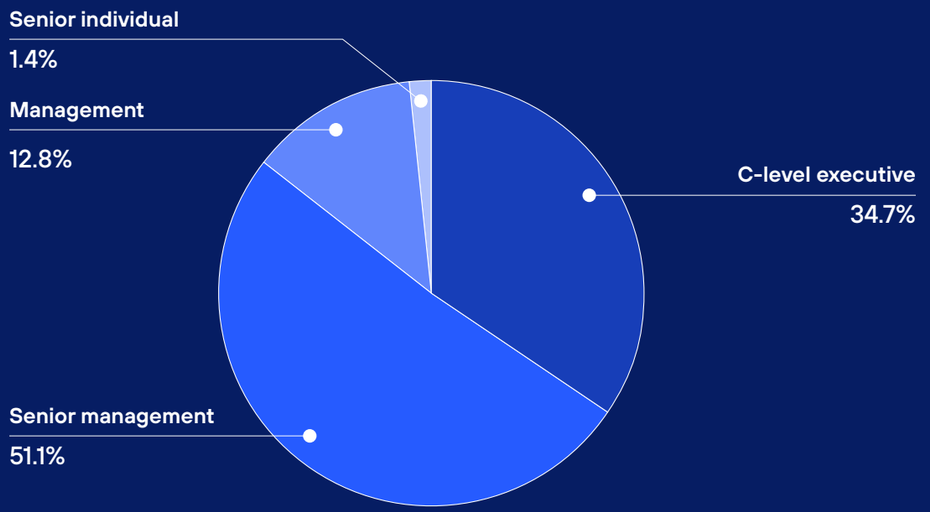

RESPONDENTS BY IDENTITY SECURITY RESPONSIBILTY

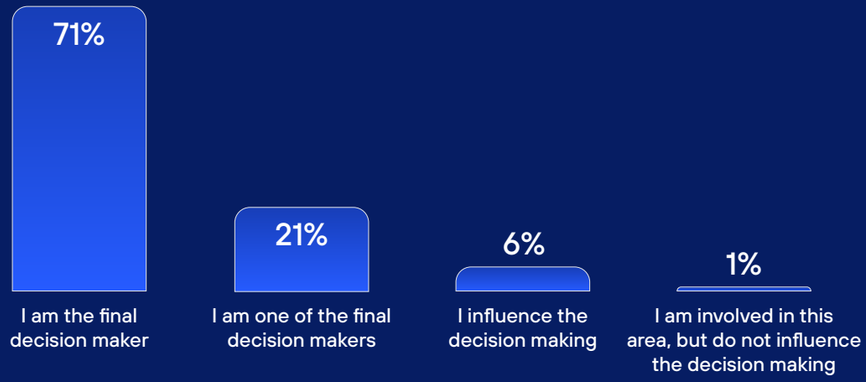

Introducing Palo Alto Networks Idira

Palo Alto Networks Idira™ is the next-generation identity security platform that secures every human, AI, and machine identity for the enterprise. Idira solutions enable organizations to discover, control, and govern the human and agentic workforce of the future. Discover risk across every identity, entitlement, and access path. Control privilege by applying zero standing privilege and just-in-time enforcement to every identity. Automate governance across the entire identity lifecycle. Explore all the ways Idira can secure the identities across your organization, visit www.paloaltonetworks.com/idira.

About Vanson Bourne

Vanson Bourne is an independent specialist in market research for the technology sector. Their reputation for robust and credible research-based analysis is founded upon rigorous research principles and their ability to seek the opinions of senior decision makers across technical and business functions, in all business sectors and all major markets. For more information, visit www.vansonbourne.com.

About Palo Alto Networks

Palo Alto Networks (NASDAQ: PANW), the global AI cybersecurity leader, protects our digital way of life with a comprehensive portfolio of cybersecurity solutions and platforms across Network, Cloud, Security Operations, AI, and Identity. Trusted by more than 70,000 customers and powered by Unit 42® threat intelligence, our AI-driven platforms eliminate complexity, empowering enterprises to modernize with confidence and securing the speed of innovation. Explore the future of security at www.paloaltonetworks.com.

3000 Tannery Way
Santa Clara, CA 95054

Main: +1.408.753.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

www.paloaltonetworks.com